

EXHIBIT B

MARKUP SPECIFICATION



METHOD FOR PERFORMING RESPONSE TIME ANALYSIS OF NETWORK

PERFORMANCE

INVENTORS

Steven J. Schaffer, Lone Tree, CO, USA

Jacob Weil, Palo Alto, CA, USA

COPYRIGHT NOTICE

[0001] A portion of the disclosure of this patent document contains material which is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure as it appears in the Patent and Trademark Office patent file or records, but otherwise reserves all copyright rights ~~whatsoever~~. The following notice applies to the software and data as described below and in the drawings hereto: Copyright © 2001, Compuware, All Rights Reserved.

FIELD OF THE INVENTION

[0002] ~~This~~The present invention relates to the field of network computing, and more particularly, to a method and system for monitoring network, client, and server performance.

BACKGROUND OF THE INVENTION

[0003] One ~~way to~~method of monitoring network performance is by measuring the processing time on a first node, such as a client, and the processing time on a second node, such as a server. ~~In earlier versions,~~ In conventional approaches, this approach was applied where the client and server were on the same LAN (local area network (LAN)), so that factors such as network delay external to the LAN did not need to be considered.

[0004] ~~To accommodate more realistic implementations,~~ network delay should also be Realistic implementations involve networks that generally include multiple LANS and interconnecting equipment and/or communications links. Thus, there is a need for an improved system and method of monitoring network performance whereby network delay is considered when monitoring network performance.

SUMMARY OF THE INVENTION

In one aspect of the invention, a method for monitoring network performance is disclosed. The method monitors a flow having one or more frames within a thread by calculating a node's active time, ~~including~~. This includes the amount of time each ~~of the frames~~frame is processed on a sending node in a network ~~and~~, the amount of time each ~~of the frames~~frame is processed on a receiving node in the network; ~~and~~ and the amount of time each ~~of the frames~~frame is in transit across the network.

BRIEF DESCRIPTION OF THE DRAWINGS

[0005] The present invention is illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings and in which like reference numerals refer to similar elements and in which:

[0006] FIG. 1 illustrates an Application Performance Report in a Chart format.

[0007] FIG. 2 illustrates an Application Performance Report in a Details format.

[0008] FIG. 3 illustrates an example of Request Preparation and Reply Preparation processing types.

[0009] FIG. 4 illustrates one exemplary embodiment.

[0010] FIG. 5 illustrates an example of a flow.

[0011] FIG. 6 illustrates an example of a multi-tier algorithm.

DETAILED DESCRIPTION OF THE INVENTION

[0012] The present invention includes various operations, which will be described below. ~~The~~ These operations ~~of the present invention~~ may be performed by hardware components or may be embodied in machine-executable instructions, which in turn may be used to cause a general-purpose or special-purpose processor or logic circuits programmed with the instructions to perform the operations. Alternatively, the operations may be performed by a combination of hardware and software.

[0013] The present invention may be provided as a computer program product ~~which~~ that may include a machine-readable medium having stored thereon instructions which may be used to program a computer (or other electronic devices) to perform a process according to the present invention. The machine-readable medium may include, but is not limited to, floppy diskettes, optical disks, CD-ROMs (Compact Disc-Read Only Memories), and magneto-optical disks, ROMs (Read Only Memories), RAMs (Random Access Memories), EPROMs (Erasable Programmable Read Only Memories), EEPROMs (Electromagnetic Erasable Programmable Read Only Memories), magnetic or optical cards, flash memory, or other type of media / machine-readable medium suitable for storing electronic instructions. Moreover, the present invention may also be downloaded as a computer program product, wherein the program may be transferred from a remote computer (e.g., a server) to a requesting computer (e.g., a client) by way of data signals embodied in a carrier wave or other propagation medium via a communication link (e.g., a modem or network connection). Accordingly, herein, a carrier wave shall be regarded as comprising a machine-readable medium.

Introduction

[0014] Response Time Analysis (hereinafter RTA) first produces underlying data ~~that is~~. This data is then formatted and presented to a user in several reports, including the Application Performance Report (Chart and Details), Node Processing Time Report, and Flows Report. ~~The RTA algorithms and techniques are the underlying technology that produces information that is~~ reports are presented in the GUI ~~(form of a~~ Graphical User Interface (GUI).

[0015] The user's first exposure to the response time analysis is in the high-level summaries presented in the health report. ~~These~~ The summaries ~~give~~ enable the user ~~quick~~ to quickly obtain critical information without ~~pouring through the~~ needing to process details. ~~Many of the details~~ Details are then made available in the Node Processing Time report and Flows ~~r~~Report.

[0016] ~~This document contains~~ The following sections give a thorough description of the algorithms and techniques used in the underlying response time analyzer RTA, as well as examples of the GUI.

Application Performance Report

Chart

[0017] FIG. 1 illustrates an Application Performance Report. The four summary results shown in the first panel of the application performance report are: given in the following paragraphs.

Network Busy Time

[0018] The Network Busy Time 100 is the total time that one or more meaningful frames are in transit across the network. Network Busy Time is computed and reported for both network directions (from primary to secondary and vice-versa). After the Flow concept has been introduced the Network Busy Time will be revisited to describe which frames are meaningful. ~~Not~~ Because only a subset of all frames generated traverse the network. ~~Only~~ frames that are exchanged between the two capture points are included in the Network Busy Time. Consequently, Network Busy Time is applicable only with a multi-segment merged or adjusted trace. A trace is a sequence of frames that have flowed between two or more nodes over the timecapture period ~~of a capture~~ (to be discussed below). Traces can be merged and adjusted, ~~for example.~~ The Network Busy Time can be broken down into ~~two parts:~~ insertion time and queuing propagation and processing time.

[0019] **Insertion Time (sec):** ~~The~~ This is the cumulative time ~~it took for~~ the frames to be inserted into the network. In a merged or adjusted trace, the insertion time for each frame is computed as $\text{AdjustedBytes} * 8 / \text{Bandwidth}$. The network bandwidth utilization is computed for each direction of the network. Only frames that are known to have traversed the network are included. The term AdjustedBytes is used in the above expression to indicate the bytes that would have ~~occurred at the point the frame crossed~~ traversed the WAN link. The capture environment allows the user to specify ~~if whether the frame size of frames should be adjusted to compensate for the different WAN headers.~~ if whether the frame size of frames should be adjusted to compensate for the different WAN headers. If the user chooses to adjust the frames do so, then

AdjustedBytes = Bytes(as captured) – DLCHeader(as captured) + DLCHeader
(specified in capture environment).

[0020] Network Queueing, Propagation & Processing (QPP) Time

(sec): ~~The portion of~~ This is the time that frames were present in the network ~~that is~~
~~caused by queueing due to queueing~~ (in routers), processing and propagation. This
~~is represents~~ the portion of the total transit time that is not counted as insertion time.
Throughout the description, this term is referred to as QPP time.

Node Active Time

[0021] ~~There is one bar for each node~~ Referring again to Fig. 1, in the
Node Active Time 102, ~~showing the portion of the task duration for which the node~~
~~was area 102,~~ there is one bar for each node that shows the active, ~~either~~ (processing
or sending) duration. Each bar is broken down into the following two components:

[0022] Node Processing Time: ~~For each node, the overall node task~~
processing time is a ~~measure of the amount of time that the node was processing~~
~~during task shown.~~ A task is a user-invoked operation that creates network traffic in the
~~course of its~~ during execution and ~~ends with~~ causes a screen update or other
acknowledgement on the user's machine. ~~A task, once~~ Once invoked, a task executes
to completion without further user intervention ~~until it has finished processing.~~

[0023] For single ~~threaded~~ applications, this is simply the sum of the
individual node processing times (described ~~later~~ below). If an application can have
multiple server requests outstanding ~~to a server~~ (such as the typical Web browser), then
the node is considered processing if it is ~~processing~~ handling one or more requests. The

Node Processing Time for any node cannot be greater than the duration of exceed the task duration, whereas the sum of the Processing Times for all nodes can be larger than exceed the task's duration if there are parallel (overlapping) threads.

[0024] Node Sending Time:— For each node, the overall node sending time is the amount of time that represents the period during which the node is in the process of sending data, but is not otherwise processing. If a node is in the process of sending a set of frames, then the node is considered to be in the sending state, but only if it is not processing another request at the same time. Sending time is important because it could indicate that the node is processing in order to prepare the remaining frames, or it could be caused by other factors that do not indicate that the node was otherwise not processing. The most likely other factor is that the network is heavily utilized and the node cannot send all of its data at once. in one transaction. Other causes potential scenarios include an insufficient TCP window size, the normal slow-start nature of TCP, inability of the receiving node to remove the data from the TCP buffer fast quickly enough, or an inefficient TCP implementation at either the sending or receiving node.

[0025] ~~From the Application Performance Chart, double clicks will result~~ Referring again to FIG. 1, the following detailed reports are available by double-clicking certain areas as discussed in the following drill-downs: paragraphs.

[0026] On a Node Processing Time portion of a node bar: Brings up the Node Processing Detail report filtered / highlighted on the specified node.

[0027] On a Node Sending Time portion of a node bar: Brings up the Node Sending Detail report filtered / highlighted on messages sent by the specified node.

[0028] On either network bar 100: Brings up the Network Utilization and Latency graph.

Details

[0029] An example of an Application Performance Report — Details is shown in FIG. 2. The major sections of the detailed report are described in the following sections.

Overall Summary

[0030] ~~The purpose of the overall summary 200 is to give the user~~
information on the task duration, any errors that were detected, the capture environment, and a ~~very terse~~ summary of the conversations, threads and turns. The information in the overall summary can alert the user to errors, ~~or the fact that they may not have captured what they intended to capture (for example, if the task time, number of conversations, or number of threads isn't what the user expects)~~ to the fact that other than the desired information was captured. The values displayed in this section are as follows.

[0031] ~~The values displayed in this section are:~~

[0031] ~~**[0032]**~~ **Task Time (sec):** ~~The~~This is the duration of the task. ~~This,~~
and should be equivalent to the task "stopwatch time," of the task. ~~If it~~this is not the
case, then portions of the time may be missing or may need to be deleted.

[0032] ~~**[0033]**~~ **Traffic Duration (sec):** ~~The span of time during~~ This is the
duration within which frames exist. The user can click on the label to open the bounce
diagram, which will show ~~them all of the~~constituent frames and ~~the time period that they~~
~~span~~their durations.

[0033] ~~**[0034]**~~ **Errors:** ~~A~~Errors A click on the Errors label opens the
Error Report, which is a graphical summary of the number and types of errors, and
warnings, ~~and informational errors~~ detected in the trace. ~~A click on the Errors label~~
~~opens the Error Report.~~

[0034] ~~**[0035]**~~ **Capture Environment:** ~~This is a legend, so to speak, for~~
indicates the meaning of the arrows ~~that are shown~~ in many other places ~~in~~within the
report. A capture is a process whereby a traffic collector or ~~capture agent~~ listens for
~~and collects~~ network frames exchanged between nodes in a distributed application and
stores ~~that~~the data for off-line analysis. The legend keys are taken from the Capture
Environment as specified by ~~a~~the user. The primary capture location is indicated on the
left and the secondary location is on the right.

[0035] ~~**[0036]**~~ **Conv:** ~~indicates~~ Conversations, both the ~~total~~ for the task
and the total, and those that ~~are occur~~ between ~~one a~~a node on the primary location and
another node in the secondary location, as indicated by the "<-->" label. Hereafter, the

phrase "conversations that traverse the network" means conversations for which one node is in the primary location and the other is in the secondary location.

[0036] ~~**[0037]**~~ **Threads:** ~~The~~ This shows both the total number of threads in the task (~~Total~~) ~~and~~ as well as in the conversations that traverse the network between the primary and secondary capture points. A thread is a sequence of frames exchanged between two nodes that constitutes a single application or protocol action. For example, the retrieval of a graphic from a WWW (World Wide Web) server is a thread. ~~A thread will always be between a pair of nodes.~~

[0037] ~~**[0038]**~~ **Turns:** ~~The~~ This is the number of turns in the task, and in the conversations that traverse the network between the primary and secondary capture points. <--> Turns specifies the sum of the turns for threads ~~that have~~ corresponding to one node in the primary location and ~~the~~ another other node in the secondary location.

[0038] ~~**[0039]**~~ **Bytes/Turn:** ~~The~~ This is the average number of bytes in each per turn, both as a total for the task and for the conversations that traverse the network.

[0039] ~~**[0040]**~~ **Frames/Turn:** ~~While (not shown, this characteristic can also be specified. It defines)~~ This is the average number of frames in each turn, both as a total for the task and for the conversations that traverse the network.

Traffic

[0040] ~~**[0041]**~~ ~~The traffic~~ Traffic section 202 provides the user with an overall a summary of several traffic measures, for the entire task (~~the~~ Total row), over the network (~~the~~ <--> row), ~~and in each direction across the network. As used herein,~~

~~“over the network” will mean, i.e.,~~ in both directions between the primary and secondary capture points), and in each direction across the network. Columns in this section ~~include:~~ are discussed in the following paragraphs.

[0041] ~~[0042] Bytes: Sum~~ This is the sum of the bytes for all frames in each classification. For the network classifications, the byte counts for each frame will be adjusted for the DLC header size if the user has chosen to ~~perform this~~ adjustment do so in the capture environment.

[0042] ~~[0043] For <-->, → and ← Bytes, the value is the number of bytes~~ that crossed the point of contention in the network as specified for the task in the Capture Environment ~~for the task.~~ If the user did not choose to adjust frames for DLC header in the capture environment, then Network Bytes = Bytes for each frame. If the user did so ~~choose to adjust frames for DLC header in the capture environment~~, then Network Bytes = Bytes – DLC Header(as captured) + DLCHeaderbytes (specified in the capture environment).

[0043] ~~[0044] % of <--> Bytes: Applicable only to the two directions of the network, this~~ This column shows ~~what the~~ percentage of the <--> bytes that traversed the network ~~are attributed to~~ in each direction. The two values will add to 100% (subject, of course, to rounding ~~in the display~~).

[0044] ~~[0045] Frames: The~~ This is the total number of frames in the task (top row), and the number of frames that should have crossed the network, whether they were contained in both captures or not. If they weren't not, such will be reported in the two Frames Missing ~~columns at~~ entries to the right of this section. A frame is a

collection of bits comprising data and control information (overhead) that is transmitted between nodes in a network.

[0045] ~~**[0046]**~~ **Avg Frame:** ~~The~~ This is the average frame size, in bytes.
~~Avg Frame =~~ , i.e., Bytes / Frames.

[0046] ~~**[0047]**~~ **Captured Load** (kbps and %): ~~The captured load~~ This is the average rate, in kbps and as a %percentage of the total bandwidth, of the frames that traversed the network in each direction. The adjusted bytes ~~(refer to the discussion, as discussed above)~~ are used.

[0047] ~~**[0048]**~~ **Frames Missing at Source:** ~~The~~ This is the number of additional frames from the sending (source) node that should have been in the trace ~~from the capture point where the sending (source) node was placed. Frames missing at the source are usually an indication that they should have been captured but were not.~~ This can be caused by the capture beginning too late or ending too early, or by the inability of the capture device to capture all of the frames.

[0048] ~~**[0049]**~~ **Frames Missing at Destination:** ~~The~~ (not shown) This is the number of additional frames that should have been ~~seen~~ in the trace from the ~~capture point where the receiving (destination) node was placed. Frames missing at the destination. Such frames could behave been lost in the~~ due to network ~~due to~~ congestion or a failure of a network component, or ~~they could be missing for the same reasons that frames missing at the source can be missing~~ reasons discussed in the preceding paragraph.

[0049] ~~**[0050]**~~ There can be other situations wherein not all of the frames that should have traversed the network were actually captured, ~~thereby resulting in missing frames.~~ For instance example, one of the captures may have started before or ended after the other and may contain frames that traversed the network. However, since those frames are not represented in the other capture, it is not known if they actually traversed the network. Such frames will be flagged with an error of Lost Frame or Dropped Frame, ~~respectively,~~ depending on whether it ~~was~~ they were captured only at ~~its~~ the source segment or destination segment, respectively.

[0050] ~~**[0051]**~~ If it is assumed that missing frames really did traverse the network, then their bytes / frames / threads / conversations are included in the <--> metrics. There is no way to know whether a frame that is missing at the destination actually consumed bandwidth at the network contention point before being dropped. Thus, the worst case is assumed —, i.e., that it ~~did consume~~ network bandwidth was consumed, and the missing frames are ~~also~~ included in the <--> metrics.

Network Busy Time

[0051] ~~**[0052]**~~ ~~The network busy time~~ The Network Busy Time section 204 helps the user determine how ~~busy~~ active the network was during the task, and how ~~that~~ busy time breaks down into insertion time and QPP time.

[0052] ~~**[0053]**~~ ~~The network busy time~~ The Network Busy Time section 204 is presented for merged tasks and single-trace adjusted tasks. There are two rows for the network metrics, one for the primary to the secondary location (→), and the second from the secondary to the primary location (←). The columns in this section correspond

exactly to the portions of the network bars 100 in the Application Performance Report chart of FIG. 1.

~~[0054] The columns in this section correspond exactly to the portions of the network bars in the Application Performance Report graph.~~

~~[0053]~~ ~~[0055]~~ **Insertion Time** (sec and % of Task Time): ~~The~~ This information represents the cumulative time that it took to insert the captured frames into the network at the point of lowest bandwidth specified by the user. ~~The insertion time is based on the network bytes of each frame should~~ Should the user decide to adjust for DLC headers of the network by the captured bytes, the insertion time is based on the network bytes of each frame. The bandwidth utilization in kbps is: determined as $(\text{Bytes that traversed the network in the specified direction} * 8) / (\text{duration of the task in seconds} * 1000)$. The bandwidth utilization in % is $(\text{the bandwidth } \underline{\text{utilization}} \text{ in kbps} / \text{capacity of the link in kbps})$, expressed as a percentage. The capacity of the link is specified by the user in the capture environment.

~~[0054]~~ ~~[0056]~~ **QPP Time** (sec and %): ~~Queueing of Task Time~~ This is the Queueing, Processing and Propagation time.

~~[0055]~~ ~~[0057]~~ **Total** (sec and % of Task Time): The total time that one or more meaningful frames was traversing the network. Meaningful frames include all data frames and TCP acknowledgements that are within the data portion of a message. Meaningful frames do not include TCP acknowledgements that are sent after the last data frame in a message is sent.

Network Frame Transit Statistics

~~[0058] The network frame transit statistics section 206 comprises 2 rows, one for each direction of the network as described above.~~

~~[0056]~~ [0059] The network frame transit statistics section 206 comprises 2 rows, one for each network direction as described above. This section includes only meaningful frames. As such, the statistics do not include TCP acknowledgements that ~~occur~~are sent after the last data frame in a message is sent.

~~[0057]~~ [0060] Transit Time (sec): ~~the first~~This column will contain~~reflect~~ the graphical depiction of the ~~min, avg,~~minimum, average and ~~max~~maximum frame transit time. The transit time of each frame is the difference between its Time Sent and Time Received.

~~[0058]~~ [0061] Transit Time Min (sec): ~~The~~This is the transit time of the frame that has the ~~smallest~~lowest transit time.

~~[0059]~~ [0062] Transit Time Avg (sec): ~~The~~This is the average (mean) transit time for the meaningful frames.

~~[0060]~~ [0063] Transit Time Max (sec): ~~The~~This is the transit time of the frame that has the largest transit time.

~~[0061]~~ [0064] Transit Time Frame Count: ~~The~~This is the number of frames included in the transit time statistics. ~~Equivalent, and is equal~~ to the number of meaningful frames that were captured at both sides (or adjusted). Note that this number is equal to or less than the number of frames that traversed the network in the

specified direction. It does not include TCP acknowledgements after a message or frames that are missing at the source or destination.

[0062] ~~**[0065]**~~ **Latency Statistics** (min, avg, max): ~~While, not shown,~~
These are statistics on the latency of meaningful frames that traverse the network ~~may be shown.~~ As described ~~before~~ above, meaningful frames are data frames and the TCP acknowledgements that occur during the data transfer portion of a flow.

[0063] ~~**[0066]**~~ **Overlap Avg (Frames):** Frame ~~The~~ This is the average number of frames that are in transit when there is at least one frame in transit. It is a measure of the application's ability to send more than one frame at a time, and the network conditions requiring the application to do so. ~~Put another way~~ In other words, it is the average number of frames sent by the application when the application has sent frames. Higher values mean the application is less susceptible to network latency and bandwidth.

[0064] ~~**[0067]**~~ **Overlap Max (Frames):** Frame ~~The~~ This is the largest number of frames that are in transit in the network at any given time.

Node Active Time

[0065] ~~**[0068]**~~ This is a tabulation ~~format~~ 208 of the node bars (processing and sending 102 (Node Processing And Sending times)) that were described earlier in the Application Performance Report Chart: of FIG. 1.

Node Processing Statistics

[0066] ~~[0069]~~ Statistics on 210 regard the node processing periods, as discussed in the following paragraphs.

[0067] ~~[0070]~~ **Processing Time (sec):** ~~A~~ This reflects a graphical depiction of the ~~min, avg~~ minimum, average and ~~max~~ maximum node processing time period.

[0068] ~~[0071]~~ **Processing Min (sec):** ~~The~~ This is the shortest node processing period.

[0069] ~~[0072]~~ **Processing Avg (sec):** ~~The~~ This is the average node processing period.

[0070] ~~[0073]~~ **Processing Max (sec):** ~~The~~ This is the longest node processing period.

[0071] ~~[0074]~~ **Processing Periods:** ~~The~~ This is the number of processing periods, and is important as it tells the user if there ~~are~~ is a large or small number of node processing periods.

[0072] ~~[0075]~~ **Overlap Avg:** ~~The~~ This is the average number of processing periods that occur simultaneously at the node during the times that the node is in at least one processing period. A value of 1.0 means that the node never processed more than one request at a time. This value cannot be smaller than 1.0.

[0073] ~~[0076]~~ **Overlap Max:** ~~The~~ This is the largest number of processing periods occurring at any given instant.

Node Processing Detail Report

[0074] ~~**[0077]**~~ In addition to the overall summary results presented in the Application Performance Report, the RTA also identifies and reports several sets of details. One of these is the node processing detail (not shown).

[0075] ~~**[0078]**~~ Each node processing time is one component in the Overall Node Processing Time. The ~~UI~~GUI allows ~~one~~the operator to see the Individual Node Processing Times for all nodes or for one node at a time. Note that since individual node processing times can overlap, the sum of the individual node processing times can be greater than the Overall Node Processing Time for a ~~node~~given node. The attributes of each node processing time, as given in the columns in the Node Processing Detail Report are listed or described in the following sections.

[0079] ~~————~~ ~~The attributes of each node processing time (aka columns in the Node Processing Detail Report) are:~~

[0076] ~~**[0080]**~~ **Node Name.**

[0077] ~~**[0084]**~~ **Node Address.**

[0078] ~~**[0082]**~~ **Errors.** ~~The~~This is the number of errors associated with either the start or the end frame. The user can ~~drill into to~~click on the error report to see the individual errors.

[0079] ~~**[0083]**~~ **Duration** ~~(sorted descending by default).~~ ~~The~~ This is the time span of the processing time, in seconds.

[0084] ~~————~~ **Processing Type.** ~~One of the list specified below~~

[0080] ~~[0085]~~ **Start Time.** ~~The~~ This is the time at which the node began processing

[0081] ~~[0086]~~ **Start Frame.** ~~The~~ This is the number of the frame number captured at the beginning of the processing time. The user can ~~drill down to~~ click on this parameter to view the bounce or packet trace ~~which will take them to~~ as of the start frame.

[0082] ~~[0087]~~ **End Time.** ~~The~~ This is the time at which the node stopped processing.

[0083] ~~[0088]~~ **End Frame.** ~~The frame~~ This is the number that ~~is of the~~ frame captured at the end of the processing time. The user can ~~drill down to~~ click on this parameter to view the bounce or packet trace ~~which will take them to~~ at the end of the frame.

[0084] ~~[0089]~~ **Start Frame Description.** ~~The~~ This is the description (decoded) of the start frame.

[0085] ~~[0090]~~ **End Frame Description.** ~~The~~ This is the description (decoded) of the end frame.

[0086] ~~[0094]~~ **Node Processing Type** This is one of the types specified in the sections below. For most of the node processing types, there are corresponding processing types for client and server. ~~For example, Client Processing and Server Processing. They basically have the same meaning, but the "Client" one~~ "Client" is assigned when the node is the client in a thread ~~whereas the "Server" one is assigned when the node is the server in the thread.~~ and vice versa. Each node processing type

has an internal code number that is listed in parentheses. ~~This code number never appears that does not appear~~ in the GUI. The node processing types and their meanings are: given in the following sections.

[0087] ~~[0092]~~ **Client Before Thread-300:** ~~The~~ This is shown in FIG. 3 at 300, and represents the time period prior to sending of the first data frame in the thread when that first data frame is sent by the client. The processing time period extends back to the previous data frame that was received by the node, or to the beginning of the task if there isn't one is no such frame.

[0088] ~~[0093]~~ **Client Processing:** ~~Within~~ This the period within a thread, when from receipt of a data frame by the client node to the time that node sends out a subsequent request, the time before that request is considered to be a client processing time. The time extends back to the previous data frame that was received by the client. Note that that. The data frame can be on the same thread or it can be on another thread.

[0089] ~~[0094]~~ **After last frame 308:** ~~The~~ This is the time period from the last data frame to the end of the task, and is always assigned to the client node for the task. You can reassign the The client node can be reassigned in the conversation map. This processing type will always end at the time that is the end of the task.

[0090] ~~[0095]~~ **Server Before Thread:** ~~This is defined in the same way~~ assimilar to Client Before Thread, but occurs if arises when the first data frame in the thread is sent by the server of the thread. Normally one would not expect the server of the thread to the server does not send the first data frame, (since the client normally

initiates activity by sending a request) but, However, if the capture starts in the middle of a thread, or if the client and server are assigned incorrectly then you may get this processing type. Note that the Thread Analysis window comprises a command to swap the client and server of a thread if they are identified incorrectly, then this processing type results.

[0091] Note that the Thread Analysis window comprises a command to swap the client and server of a thread if they are identified incorrectly.

[0092] ~~[0096]~~ **Server Processing 304:** Within a thread, this This is the time period within a thread from the point that receipt of the last frame in a request is received by the server to the time that the first response frame is returned by the server. During this time, the server is assumed to be processing the request and consequently this time is identified as a server processing time. Note that with multi-tier applications (discussed below) there are cases where an upper-tier request may interrupt a server processing time. Multi-tier is described later.

[0093] ~~[0097]~~ **Request Preparation 302:** In This processing type arises in multi-tier, when applications. It is the period from receipt of a request by a mid-level server processes after receiving a request (from a lower tier) until it the mid-level server begins sending a subsequent request to another server. ~~FIG. 3 shows this processing type at the App Server between frames 1 and 3.~~

[0094] ~~[0098]~~ **Reply Preparation 306:** In This processing type also arises in multi-tier, when applications. It is the period from receipt of a reply by a mid-level server processes after receiving a reply (from another server until it begins

~~sending its) to initiation of a reply to its~~the requesting node. ~~FIG. 3 shows this processing type at the App Server between frames 5 and 7.~~

Flows Report

[0095] ~~[0099]~~ A flow is a set of data frames that is sent from one node to another. ~~Specifically, a flow,~~ comprises only frames in a single thread, and spans a time period where there areduring which no data frames traveling in the opposite direction. A flow ~~also includes the TCP acknowledgements that are sent in the opposite direction during the flow and before the direction of data transmission~~ direction reverses.

[0096] ~~[0100]~~ Meaningful frames ~~were,~~ discussed earlier. All above, comprise all data frames, and TCP acknowledgements within a data flow ~~are meaningful frames.~~ TCP acknowledgements that occur after the last data frame in a flow are not meaningful frames for the purpose of network busy time computation.

[0097] ~~[0101]~~ In the flows report, each flow ~~has the following~~includes a number of attributes ~~(, arranged in columns),~~ as discussed in the following sections.

[0098] ~~[0102]~~ Errors. A The flows report provides a graphical depiction of the ~~number of~~relevant errors and warnings, ~~as is done in~~with other reports. ~~Includes a~~ A link to the error report ~~where the errors belonging to frames in the flow would be shown.~~ is provided. Since a flow comprises one or more frames, the errors identified by or associated ~~to~~with any frame in the flow should be included in this summary.

[0099] ~~[0103]~~ Sending Node ~~(This is the name and address).~~ The of the node that ~~sends~~the data frames in the flow. This node will ~~also~~ receive TCP

acknowledgements from the corresponding receiving node.

[0100] ~~{0104} Receiving Node~~ (This is the name and address). ~~The~~ of
the node that receives d the data frames in the flow. This node will ~~also~~ send TCP
acknowledgements to the corresponding sending node.

[0101] ~~{0105} Data Duration~~ (~~seconds~~): ~~The time~~ This is the period in
seconds from ~~when~~ the time the sending node sent the first frame in the flow to the time
that the receiving node received the last data frame in the flow. This is related to other
fields ~~as follows:~~ by the relationship $\text{Data Duration} = (\text{End Data Time} - \text{Start Time})$

[0102] ~~{0106} Avg Data Rate~~ (~~bits/sec~~). ~~The~~ This is the average data
rate in bits/sec during the flow. This ~~is~~ information may be important to the user,
because flows with low data rate may be ~~worthy of further~~ demand investigation. For
example, the user may need to determine why the data is not being transferred more
quickly, particularly if the flow is also longer than most other flows. ~~Computed~~ This is
computed as $(\text{Data Payload Bytes} * 8) / (\text{End Data Time} - \text{Start Time})$.

[0103] ~~{0107} Bytes~~. ~~The~~ This is the total number of bytes in all ~~of the~~
frames in the flow.

[0104] ~~{0108} Data Payload Bytes~~. ~~The~~ This is the sum of the payload
bytes ~~for in~~ all ~~of the~~ frames in the flow.

[0105] ~~{0109} Frames~~. ~~Number~~ This is the number of frames in the flow.

[0106] ~~{0110} Data Frames~~. ~~Number~~ This is the number of data frames
in the flow.

[0107] ~~{0111}~~ **First Frame.** ~~The~~ This is the sequence number of the first frame in the flow.

[0108] ~~{0112}~~ **Last Data Frame.** ~~The~~ This is the sequence number of the last data frame in the flow.

[0109] ~~{0113}~~ **Last Frame.** ~~The~~ This is the sequence number of the last frame, either data or acknowledgement. If there are TCP acknowledgement frames after the last data frame, then this will differ from the Last Data Frame.

[0110] ~~{0114}~~ **Start Time.** ~~The~~ This is the time that the first data frame was sent.

[0111] ~~{0115}~~ **End Data Time.** ~~The~~ This is the time that the last data frame was received.

[0112] ~~{0116}~~ **End Time.** ~~The~~ This is the time that the last frame (data or acknowledgement) was received. Note that this is normally not important because trailing TCP acknowledgements do not have an impact on the response time.

[0113] ~~{0117}~~ **Other columns that may be included are:**

[0114] ~~{0118}~~ **Data Direction.** ~~Either~~ This reflects primary-to-secondary direction or vice versa. ~~Best, as~~ indicated with arrows (→ and ←) arrows. For flows that are within a capture location, ~~should read~~ the caption "within <capture point>" appears, where <capture point> indicates the location of interest.

[0115] ~~{0119}~~ **Network Busy Time (seconds).** ~~The~~ This is the total time in seconds that one or more frames was in transit during the flow.

[0116] RTA Algorithm Details

[0117] *Overall Approach*

[0118] ~~**[0120]**~~ The RTA functions first at the thread level. Each thread is assumed to be a single-threaded sequence of request / response exchanges between the client of the thread and the server of the thread. It is the responsibility of the protocol decoder to ensure this requirement. A thread is broken down into 5 time periods that fit into, described in the following 5 categories: sections.

[0119] ~~**[0121]**~~ Client Processing Time.

[0120] ~~**[0122]**~~ Client Sending (Flow is being sent from client to server).

[0121] ~~**[0123]**~~ Server Processing Time.

[0122] ~~**[0124]**~~ Server Sending (Flow is being sent from server to client)

[0123] ~~**[0125]**~~ Processing interrupted by another thread (This period only occurs in the case of multi-tier or overlapping requests at the client)

[0124] *Exemplary Embodiment - Single Thread*

[0125] ~~**[0126]**~~ RTA concepts can be illustrated in accordance to one exemplary embodiment as shown in FIG. 4. FIG. 4 ~~shows~~ is a bounce diagram for a typical client / server or Web application. (The application could be anything, e.g., a 2-tier SQL application, a web browser to a web server, or even an application using home-grown based on ad hoc protocols.)

[0126] ~~**[0127]**~~ Assume that all of these frames are in shown exist within a

single thread. In this example, the client sends a 2-frame request to the server, the server processes ~~for over~~ a bit period, and then the server sends a 3-frame reply back to the client. The diagram shows the data frames that would be exchanged, as well as exemplary TCP acknowledgements ~~that will be seen if the application uses TCP/IP.~~ Note that TCP is a very dynamic protocol, and therefore may not send a TCP acknowledgement for every frame. ~~So Accordingly, the diagram below illustrates only one of the many that could have occurred.)~~ many possible variations.

[0127] Node Processing and Sending

[0128] ~~[0128]~~ The Application Performance Report — Chart of FIG. 1 would show ~~that the processing time for the client was as~~ the sum of the two processing times identified in the figure — ~~the~~ FIG. 4, i.e., one at the beginning and ~~the~~ one at the end.

[0129] ~~[0129]~~ The processing time for the server is just the single processing time ~~in the middle of the trace.~~

[0130] ~~[0130]~~ The 304, and the sending time for both nodes is indicated in the figure. as indicated in the FIG. 4.

[0131] Flows

[0132] ~~[0131]~~ In this example there are two flows. The first flow is sent from the client to the server and comprises frames 1 through 4. The last data frame in this flow is frame 3. The overall flow duration and flow data duration ~~of the first flow~~ is are annotated in FIG. 5. Because ~~f~~Frame 4 is a TCP acknowledgement that is sent after the last data frame (in the flow (Frame 3) is sent ~~in the flow~~, it is not considered a

meaningful frame and Thus, the network is not considered busy for the time that frame 4 is in transit.

[0133] ~~**[0132]**~~ A similar Similar analysis could be made ~~for~~ applies to the second flow in this example. The second flow is sent from the server to the client, and comprises frames 5 through 10. The data duration ~~is from~~ for this flow spans the time frame 5 is sent to the time frame 8 is received.

[0134] Network Frame in Transit and Latency Statistics

[0135] ~~**[0133]**~~ Again ~~referencing~~ referring to FIG. 4, the times when a frame is in transit can be seen. ~~The~~ In general, the TCP acknowledgements that are returned after a ~~flow~~ data is completely received have no impact on the overall response time. ~~;~~ Therefore, they are omitted from the network Frame in Transit measure and Latency statistics. ~~In~~ Accordingly, in this example frames 4 and 10 do not impact the user-perceived application response time, and thus they are not included in the network latency measures.

[0136] ~~**[0134]**~~ A high network frame in transit time can be an indication that the combination of network latency and the application's sequential request / response behavior is affecting the response time. ~~If you see a~~ A high network busy time, ~~you should also look at the network utilization to see if the cause of the high busy time is~~ may be caused by insufficient network bandwidth. ~~This can be easily done~~ This may be investigated by consulting the network utilization information in the Performance Report — Chart ~~since the~~ This chart breaks down a network frame in transit ~~is broken down into components caused by bandwidth and latency. If the bandwidth portion of~~

~~the bar~~component dominates, then ~~the~~ lack of bandwidth is causing the network to be busy.

[0137] *Exemplary Embodiment - Multi-Thread*

[0138] Node Processing and Sending Time

[0139] ~~[0135]~~ Node processing time analysis becomes more complicated when the application is multi-threaded. In the example above, the sum of the node processing times equals the node active times. When ~~there are overlapping node~~ processing times overlap at a node, only one of such them is counted and; consequently, the sum of the individual node processing times can be greater than the calculated overall node processing time.

[0140] ~~[0136]~~ The user is shown results for both node processing and sending times. ~~Node processing is when the node is definitely processing. Node sending times mean that the node might be processing and consequently this time might contribute to the total response time. Sending times could be caused by other~~ Node processing time unambiguously reflects periods during which the node is processing requests (as a server) or processing a reply prior to sending the next request (as a client). Conversely, node sending time reflects not only node processing, but potentially other activity as well. It is difficult to determine what contributed to the node sending time without examining the individual flows sent. For example, node sending times could result from factors such as the receiving node's inability to process incoming data ~~fast~~quickly enough, insufficient bandwidth in the network, the sending node's inability to make all of the data available quickly enough, or a TCP window size

that is too small. ~~One can drill down from the node sending~~ Consequently, node sending time might contribute to the total response time. For example, in cases where the network is heavily utilized or has high latency, a high node sending time can be caused by limitations of the network. To explore this, a user can link from the node-sending-time bar to view the flows sent by that node. From there, the user should look at the flows and The user could thereby attempt to identify if whether the time lapse between frames sent by the node was really represented actual processing time or not. Often this determination will require some knowledge of the application. ~~In cases where the network is heavily utilized or has high latency, a high node sending time can be caused by the network.~~

[0141] ~~[0137]~~ Node Processing Time can tell you how much of the time the node was processing requests (as a server) or processing a reply prior to sending the next request (as a client). ~~The aggregate sending time is important because it can tell you how much of the time the node was trying to send data to other nodes. It is difficult to draw absolute conclusions about what caused the sending time without looking at the individual flows that the node sent.~~

[0142] Node Processing Times

[0143] ~~[0138]~~ The Node Processing Detail report (not shown) shows the individual processing times that were detected for each node. The report is sorted initially in arranged in order of descending duration so that you can easily see, with the largest individual processing times at the top. ~~Node processing times occur at~~ Processing time at a client node node begins just prior to the node sending out the first request in a thread, and then ends within a thread prior to each succeeding request that

the node sends. Processing times occur at servers from the time at a server node begins when the server receives a response until the time request and ends when that it server begins sending a reply.

[0144] ~~**[0139]**~~ For further understanding of node processing times, a Node Processing Detail report can be opened on a trace. The window can then be split, and the packet trace placed at the bottom. For each node processing time ~~that is~~ selected, there will be two frames surrounding the processing time. For client processing times, there will be a prior reply (or the beginning of the trace) and the request that the client sends at the end of its processing time. For server processing times, there will be the request (actually the last frame in the request if it is a multi-frame request) and the response (actually the first frame in the response if it is a multi-frame response).

[0145] Flows

[0146] ~~**[0140]**~~ As discussed previously, there can be many causes for a high node sending time, and the cause of this can be difficult to determine.

[0147] Overlapping Threads

[0148] ~~**[0141]**~~ Any time two or more threads overlap, more than one node can be processing or sending at a time. ~~Or~~ Alternately, a single node could be processing or sending on more than one thread at the same time. These processing and sending times are aggregated by concluding that:

[0149] ~~**[0142]**~~ A node is processing if it is processing on one or more threads,

[0150] ~~[0143]~~ A or that a node is sending if it is sending on one or more threads and is not processing.

[0151] *Exemplary Embodiment - Multi-Tier*

[0152] ~~[0144]~~ FIG. 6 may be used to describe FIG. 6 describes the handling of multi-tier. ~~The multi-tier algorithm works as follows:~~

[0153] ~~[0145]~~ At each point in applications, as follows. At each time that a data frame enters a node, the time, frame seq# and its thread is recorded in member variables of the CNode class.

[0154] ~~[0146]~~ When a client of a node sends out the first frame in a request, there will be a processing time. To determine its type, it is determined if whether the most recent data frame that arrived at the node was a request frame from another client. If it was, then ~~<MISSING TEXT>~~ the type is 'Request Preparation,' otherwise, the type is 'Client Before Thread.'

[0155] Conclusion

[0156] ~~[0147]~~ Thus, an exemplary A method and system for performing response time analysis of network performance have been described. Although the present invention has been described with reference to specific exemplary embodiments, it will be evident that various modifications and changes may be made to these embodiments without departing from the broader spirit and scope of the invention. Accordingly, the above description and drawings are to be regarded in an illustrative rather than a restrictive sense.

WHAT IS CLAIMED IS:

1. A method to measure an application's performance in a network, comprising within a thread, monitoring a flow having one or more frames by calculating:

an amount of time each of the frames is processed on a sending node in a network;

an amount of time each of the frames is processed on a receiving node in the network; and

an amount of time each of the frames is in transit across the network.
2. A method to analyze network performance resulting from a task, comprising:

displaying a first time representing a time that one or more meaningful frames are in a network traveling in a first direction;

displaying a second time representing a time that one or more meaningful frames are in the network traveling in a second direction;

displaying a third set of times representing times that each of one or more nodes in the network is active.
3. The method of claim 2, wherein the first and second times, and the third set of times are displayed in a chart, and:

each of the first and second times is represented by a bar that shows at least one of:

a cumulative time that it took for the one or more meaningful frames to be inserted into the network;

the portion of time that the one or more meaningful frames were in the network as a result of queueing in routers, processing, and propagation (QPP).

the third set of times for a given node is represented by a bar that shows at least one of:

a total amount of time that the given node was processing during the task;
and

a total amount of time that the given node was sending during the task, but not processing.

4. The method of claim 3, wherein the insertion time for each frame is computed as $\text{AdjustedBytes} * 8 / \text{Bandwidth}$, AdjustedBytes being used to indicate the bytes that would have occurred at the point the frame crossed a WAN (Wide Area Network) link in the network.
5. The method of claim 2, wherein the first and second times, and the third set of times are displayed in a detailed report.

6. The method of claim 5, wherein the detailed report comprises one or more of:
- an overall summary comprising the duration of the task;
 - a traffic section comprising byte and frame information;
 - a network busy time section comprising information about how busy the network was during the task, and how the busy time breaks down into insertion time and QPP time;
 - network frame transit statistics section comprising various transit times for each frame;
 - a node active time section comprising information about the processing and sending times for each node in the network; and
 - a node processing statistics section comprising statistics on the node processing periods.
7. A method to monitor network performance resulting from a task, comprising displaying a processing time corresponding to a first node in the network, each processing time having one or more attributes, including a processing type.
8. The method of claim 7, wherein the processing type comprises any one of: a time period prior to a first data frame in a thread sent by a client;
- a time period prior to a subsequent request within a thread is sent by the client;

a time period from a last data frame to the end of the task;

a time period prior to a first data frame in a thread sent by a server;

a time period from the point that the last frame within a thread in a request is received by the server to the time that a first response frame is returned by the server;

a time period that a first server processes after receiving a request from a lower tier until the first server begins sending a subsequent request to a second server;

a time period that a first server processes after receiving a reply from a second server until the second server begins sending its reply to a third server, the third server being a requesting node.

9. The method of claim 7, additionally comprising displaying one or more processing times, each processing time corresponding to a node in the network that is not the first node.

10. The method of claim 9, wherein each processing time additionally has at least one of the following attributes:

the number of errors associated with one of a start frame and an end frame;

a time span of the processing time;

a time at which the node corresponding to the processing time began processing;

a time at which the node corresponding to the processing time stopped
processing;

a start frame representing a frame number at the beginning of the processing
time;

a description of the start frame;

an end frame representing a frame number at the end of the processing time;
and

a description of the end frame.

11. A method to analyze network performance, comprising generating a flows report to monitor a given flow, the given flow having one or more frames that are sent from a sending node to a receiving node, and the flows report having one or more attributes including:

an errors attribute depicting the number of errors belonging to the one or more
frames;

a sending node attribute indicating the sending node;

a receiving node attribute indicating the receiving node;

a data duration attribute indicating a time period from when the sending node sent a first frame in the flow to the time that the receiving node received the last frame having data in the flow;

an average data rate attribute indicating an average data rate for the flow;

a bytes attribute indicating a total number of bytes in the frames in the flow;

a data payload bytes attribute indicating a sum of the payload bytes for the frames in the flow;

a frames attribute indicating a number of frames in the flow;

a data frames attribute indicating a number of frames having data in the flow;

a first frame attribute indicating a sequence number of the first frame in the flow;

a last data frame attribute indicating a sequence number of the last frame having data in the flow;

a last frame attribute indicating the sequence number of the last frame, having one of data and acknowledgement;

a start time attribute indicating a time that the first frame having data was sent;

an end data time attribute indicating a time that the last frame having data was received;

an end time attribute indicating a time that the last frame, having one of data and acknowledgement, was received;

a data direction attribute indicating a direction in which the flow was traveling;
and

a network busy time attribute indicating a total time that the one or more frames was in transit during the flow.

ABSTRACT OF THE INVENTION

A system and method and apparatus are described for analyzing the performance of a network while processing an application. The method involves measuring and monitoring network performance by calculating the amount of time nodes in a network are active processing and sending frames, as well as the amount of time that the frames spend traversing the network. Graphical user interfaces are provided to effectively present significant measurements and calculations.